
aws
INNOVATE
ONLINE CONFERENCE

aws INNOVATE
ONLINE CONFERENCE

分会场二：基础设施

计算的创新
AWS Graviton 系列揭秘

张洋，AWS 解决方案架构师经理

议程安排

AWS Graviton 系列回顾

AWS Graviton 2 系列介绍

AWS Graviton 2 深度分析

AWS M6G 系列场景性能展示

最广泛的处理器选择

Intel® Xeon Scalable processors

AMD EPYC processors

Graviton processors

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

第一代 AWS Graviton processor

基于 64位 Arm Neoverse 核心由 AWS 进行定制

Custom AWS silicon with 64-bit Arm Neoverse cores

为优化云原生负载而生

Targeted optimizations for cloud-native workloads

助力您快速创新，构建和迭代

Rapidly innovate, build, and iterate on behalf of customers

AWS Graviton processor 赋能的 EC2 A1实例

高可用性

9 个区域

北美 (N. Virginia, Oregon, Ohio)

欧洲 (Ireland, Frankfurt)

亚太 (Mumbai, Sydney, Tokyo, Singapore)

应用场景

横向扩展 负载

Web 层

容器化 微服务

基于ARM的软件应用开发

最高节省 45% 费用

灵活配置

6 个实例类型

提供 裸金属 配置

最多 16 vCPUs, 32GiB 内存

最多 10 Gbps 网络, 3.5 Gbps EBS

构建于 AWS Nitro 系统平台

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

Amazon EC2 A1 详细参数

Instance	CPU Arch	vCPUs	Memory (GB)	Network bandwidth (Gbps)	Amazon EBS-optimized	Amazon EBS bandwidth (Mbps)	EBS-optimized burst bandwidth (Mbps)
a1.medium	arm64	1	2	Up to 10	Yes	300	Up to 3500
a1.large	arm64	2	4	Up to 10	Yes	525	Up to 3500
a1.xlarge	arm64	4	8	Up to 10	Yes	800	Up to 3500
a1.2xlarge	arm64	8	16	Up to 10	Yes	1750	Up to 3500
a1.4xlarge	arm64	16	32	Up to 10	Yes	3500	3500

5 instance sizes, up to 10 Gbps networking and EBS-optimized burst
US East (N. Virginia and Ohio), US West (Oregon), Asia Pacific (Mumbai, Singapore, Sydney, and Tokyo), and Europe (Frankfurt and Ireland)
More regions coming soon!

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

A1 versus C5/C4

	a1.4xlarge	c5.9xlarge	c4.8xlarge
CPU	Alpine AL73400	Cascade Lake / Platinum 8124M	Intel Xeon E5-2666 v3
ISA	ARM v8	x86	x86
Microarchitecture	Cortex-A72	Cascade Lake / Skylake	Haswell
Core count	4 * 4	18	18
Thread count	16	18 * 2	18 * 2
Base frequency	2.3 GHz	3.0 GHz	2.9 GHz
Turbo	None	3.5 GHz (1 core) 3.4 GHz (all cores)	3.5 GHz (1 core) 3.2 GHz (all cores)
L1i cache	48 KB * 16	32 KB * 18	32 KB * 10
L1d cache	32 KB * 16	32 KB * 18	32 KB * 10
L2 cache	2 MB * 4	1 MB * 18	256 KB * 10
L3 cache	None	1.375 MB * 18	25 MB shared
Launched	2018 Q4	2017	2015

网络性能分析

	CPU & Memory	NIC Queue	Cost ($ per hour)
a1.large	2core 4GB	1	0.051
c4.large	2core 3.75GB	2	0.10
c5.large	2core 4GB	2	0.085
a1.xlarge	4core 8GB	1	0.102
c4.xlarge	4core 7.5GB	2	0.199
c5.xlarge	4core 8GB	4	0.17

- Performance: A1 ≈ C5 > C4
- Pricing: A1 < C5 < C4
- 40% off to C5, 49% off to C4

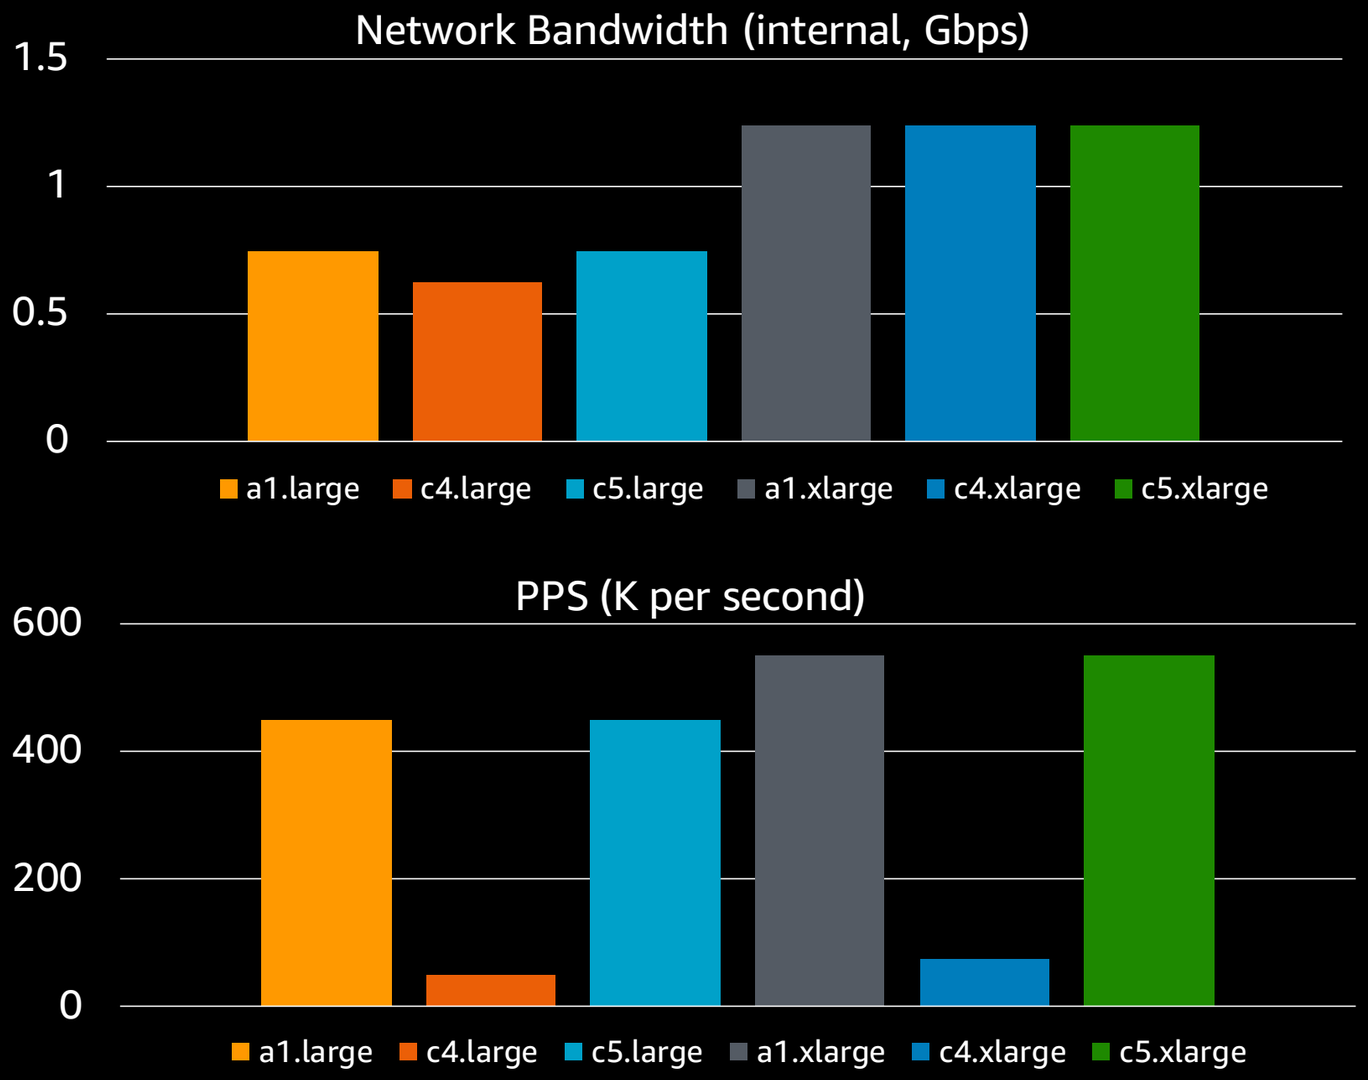

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

A1 实例非常适合网络敏感的应用类型

AWS Graviton2 Processor

AWS Graviton 2 赋能 6 系列新实例

通用型	计算优化型	内存优化型
4GB DRAM/vCPU	2GB DRAM/vCPU	8GB DRAM/vCPU
M6g	**C6g**	**R6g**
M6gd	**C6gd**	**R6gd**
可申请 M6g 预览版	2020 敬请期待	

所有实例配置增强型网络、EBS 和 3 个本地 NVMe 闪存

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

AWS Graviton2 处理器

AWS Graviton2 processor

4X	5X	7X
更多内核	更快内存	更高性能

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

AWS Graviton 2 处理器: 深入探索

aws INNOVATE
ONLINE CONFERENCE

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

AWS Graviton 2 宏观对比

AWS Graviton processor

- AWS 第一款 ARM 处理器
- 最大 16 核
- A1 系列
- 16 nm 工艺
- ~50 亿个晶体管

AWS Graviton2 processor

- 不断创新改进
- 最大 64核
- M6G系列
- 7 nm 工艺
- ~300 亿个晶体管
- 7 倍性能提升
- 每 vCPU 约 2 倍性能提升

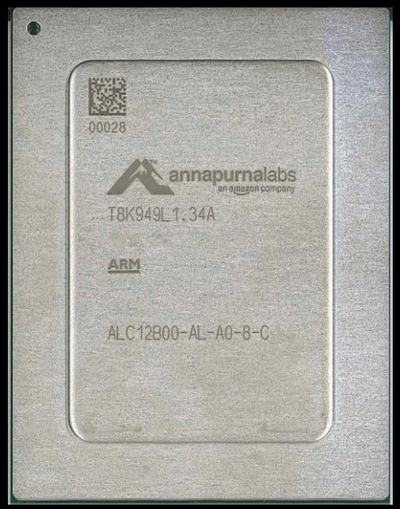

AWS Graviton 2 – 内核

- Arm Neoverse N1 内核
- Arm v8.2 指令集
- ARM N1 架构的经典实现
  - 每vCPU 64KB 1级指令与数据缓存，1MB 2级缓存
  - 指令缓存连贯性
  - 中断，虚拟化，与上下文切换的低开销
  - 4-wide front-end, with 8-wide dispatch/issue
  - 双倍 SIMD 单元
  - 使用 int8, fp16 指令加速机器学习
- 每个 vCPU 都是物理核心
  - 无超线程设计 (SMT)

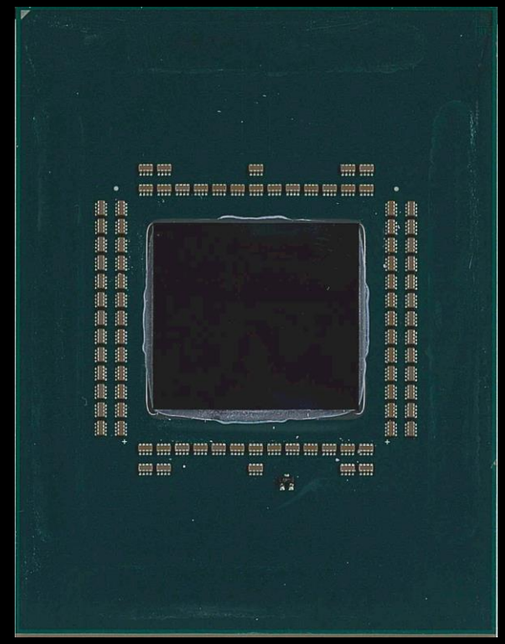

Graviton2 – 互联

- 64 个核心通过 mesh 结构互联
- ~2 TB/s 对分带宽
- 32 MB LLC 三级缓存
  - 近 100 MB 用户可使用缓存
- 无 NUMA 设计
  - 每个内核访问其它内核与内存的路径一致
- 64 通道 PCIe gen4
  - Provide flexibility for different instance configurations

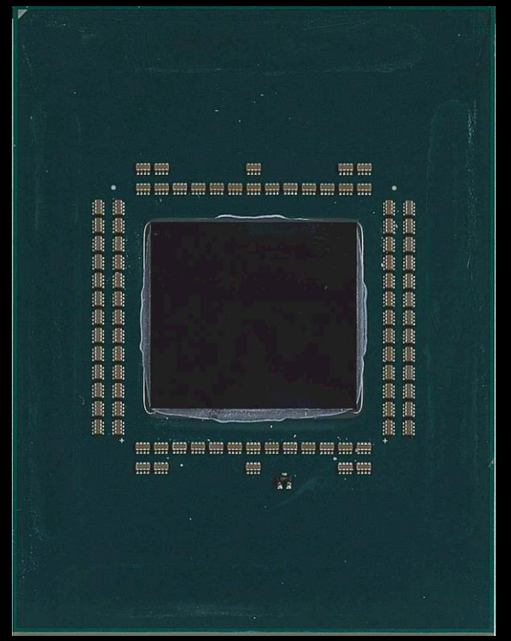

Graviton2 – 系统

- 8x DDR4-3200 通道 → 超过 200GB/s
  - 使用临时密钥与 AES-256 加密算法加密内存访问
  - 所有 CPU 核心访问内存享有一致性延迟
- 1Tbit/s 压缩算法加速器
  - 2xlarge 配置以上的实例将包含硬件级加密器
  - 发布前 DPDK 与 Linux 内核已经成功发布
  - 压缩效率高达 15GB/s 解压效率高达 11GB/s

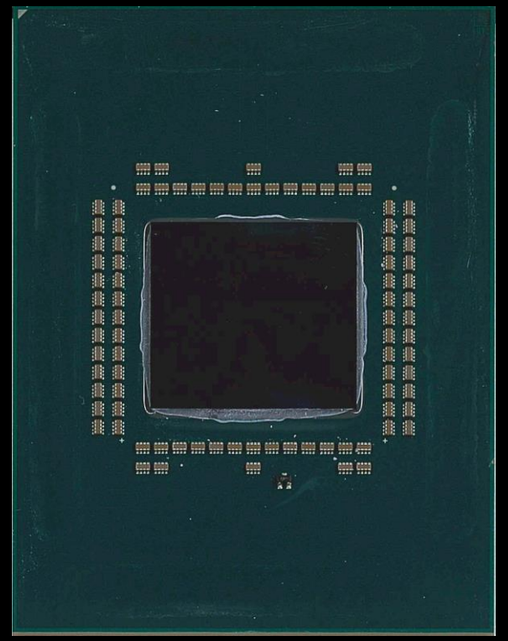

Graviton2 赋能实例

Graviton2

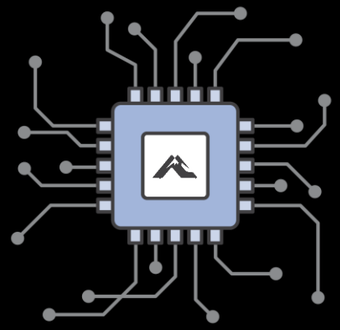

Industry leading performance

Nitro Security Chip

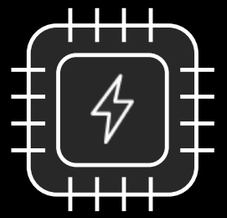

Integrated into motherboard
Protects hardware resources

Nitro Card

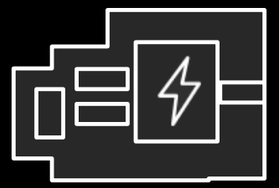

Amazon Elastic Block Store,
Elastic Network Adapter
Monitoring, and security

Nitro Hypervisor

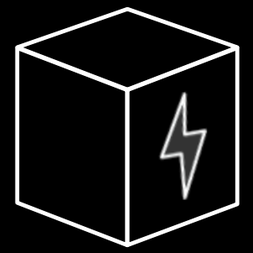

Lightweight hypervisor
Memory and CPU allocation
Bare Metal-like performance

模板化设计，专业为您打造

支持更加广泛的应用场景

Web and gaming servers

Open-source databases

High performance computing

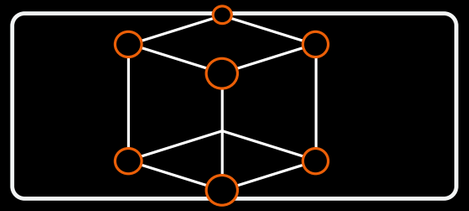

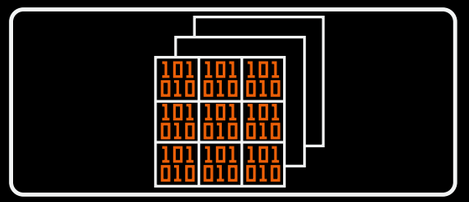

In-memory caches

Media encoding

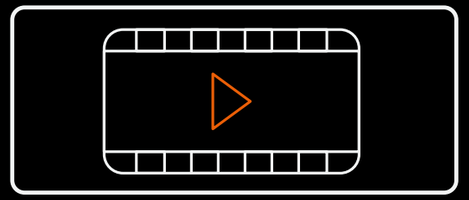

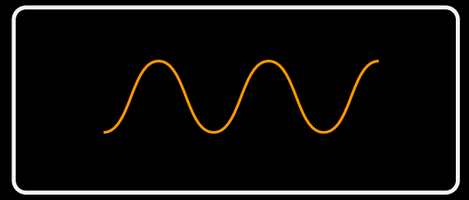

EDA

Analytics

Microservices

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

Amazon EC2 M6 具体参数表

20% lower cost vs. M5

Instance	vCPUs	Memory (GB)	Network Bandwidth (Gbps)	EBS Optimized	EBS Bandwidth (Mbps)	EBS Optimized Burst Bandwidth (Mbps)
m6g.medium	1	4	Up to 10	Yes	315	4,750
m6g.large	2	8	Up to 10	Yes	630	4,750
m6g.xlarge	4	16	Up to 10	Yes	1,188	4,750
m6g.2xlarge	8	32	Up to 10	Yes	2,375	4,750
m6g.4xlarge	16	64	Up to 10	Yes	4,750	4,750
m6g.8xlarge	32	128	12Gbps	Yes	9,000	9,000
m6g.12xlarge	48	192	20Gbps	Yes	13,500	13,500
m6g.16xlarge	64	256	25Gbps	Yes	18,000	18,000

实测性能参考

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

aws INNOVATE
ONLINE CONFERENCE

SPEC cpu2017

- Industry standard CPU intensive benchmark
- Run on all vCPUs concurrently
- Comparing performance/vCPU

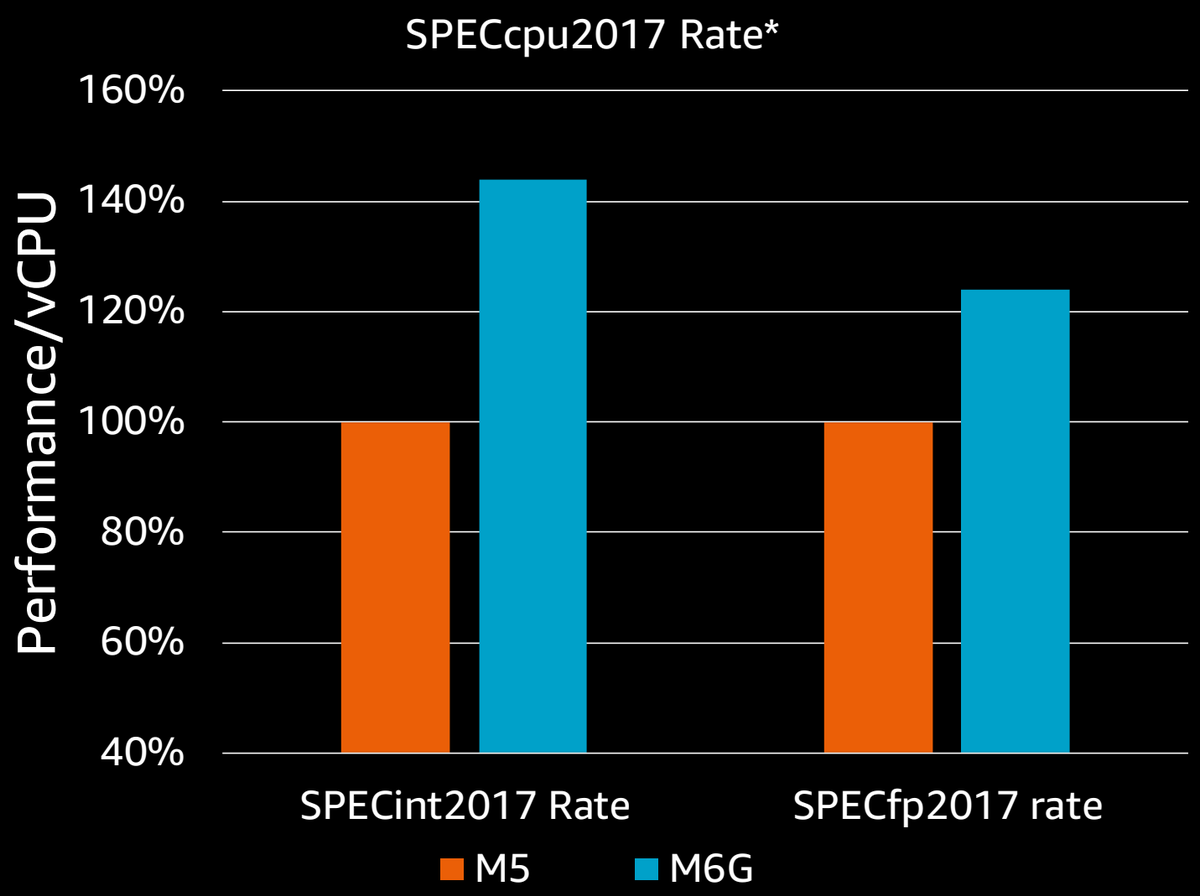

* All SPEC scores estimates, compiled with GCC9 -O3 -march=native, run on largest single socket size for each instance type tested.

Load Balancing with Nginx

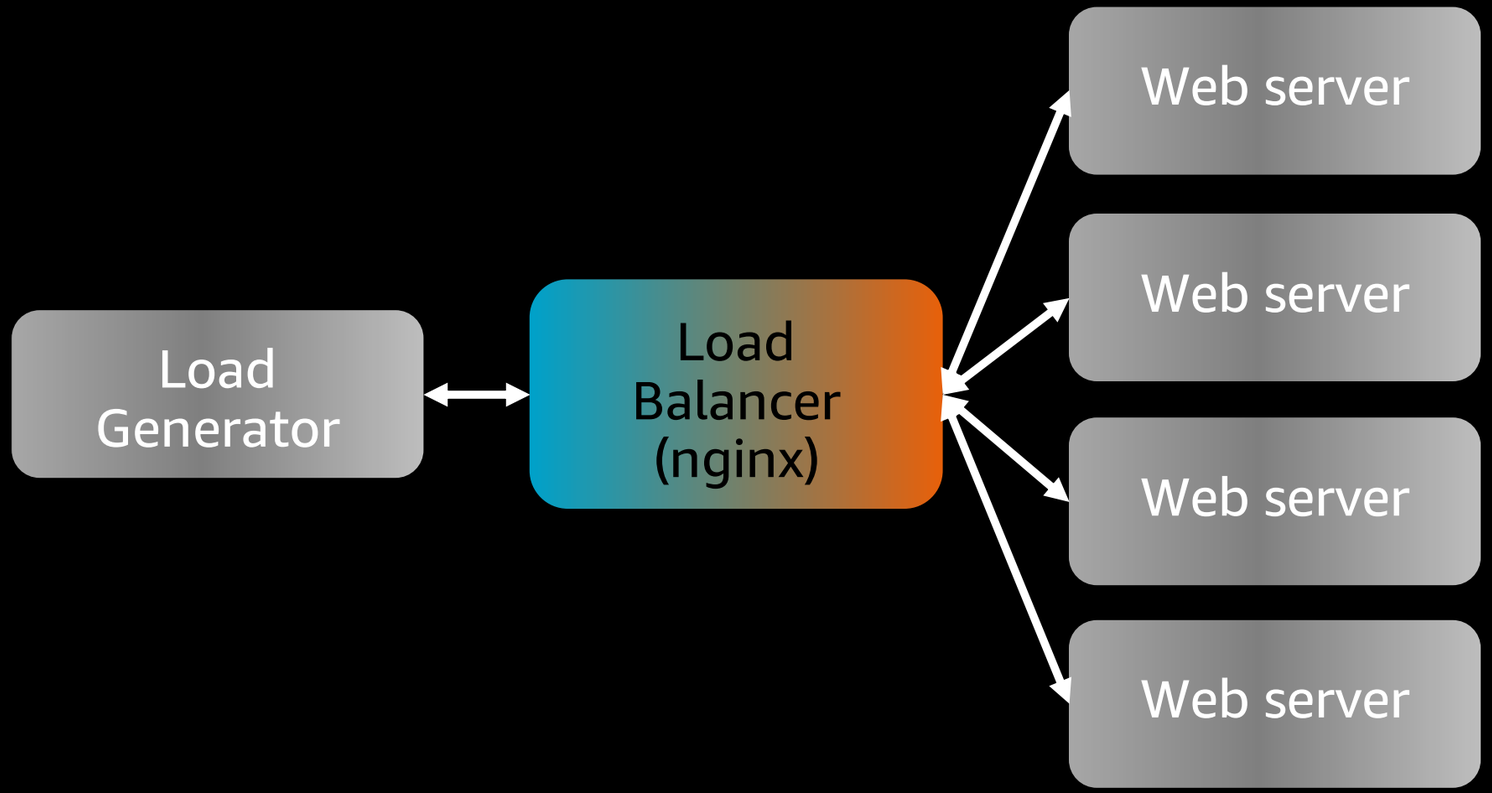

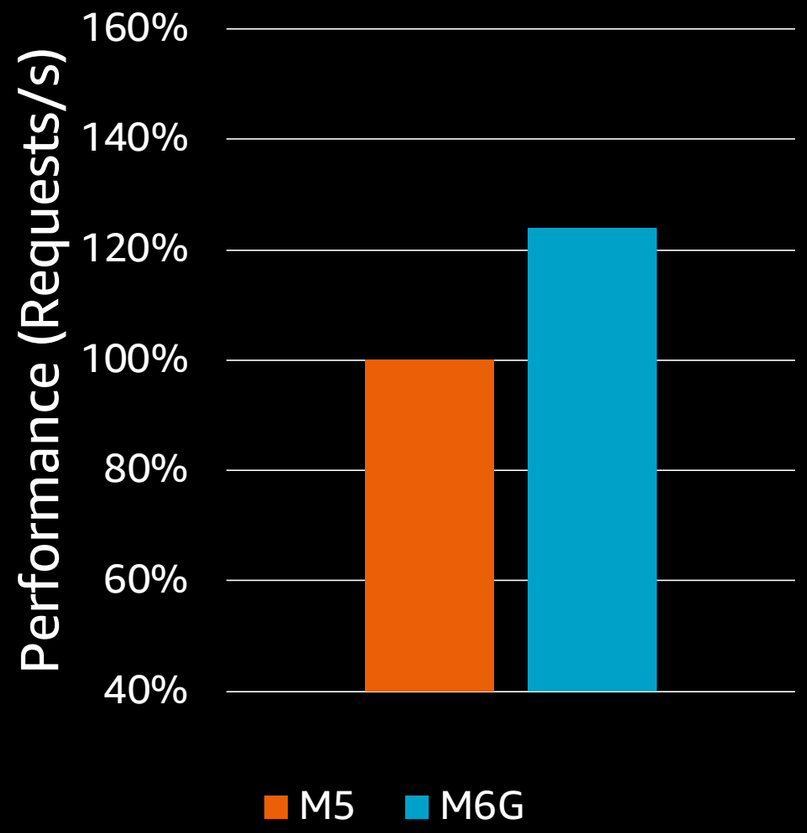

NGINX v1.15.9, 512 clients, 128 GET/POST payloads, all HTTPS connections, AES128-GCM-SHA256, OpenSSL 1.1.1, 4 target machines, all tests run on 4xl size; load generator c5.9xl; web servers c5.4xls; All servers run in a cluster placement group

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

Memcached

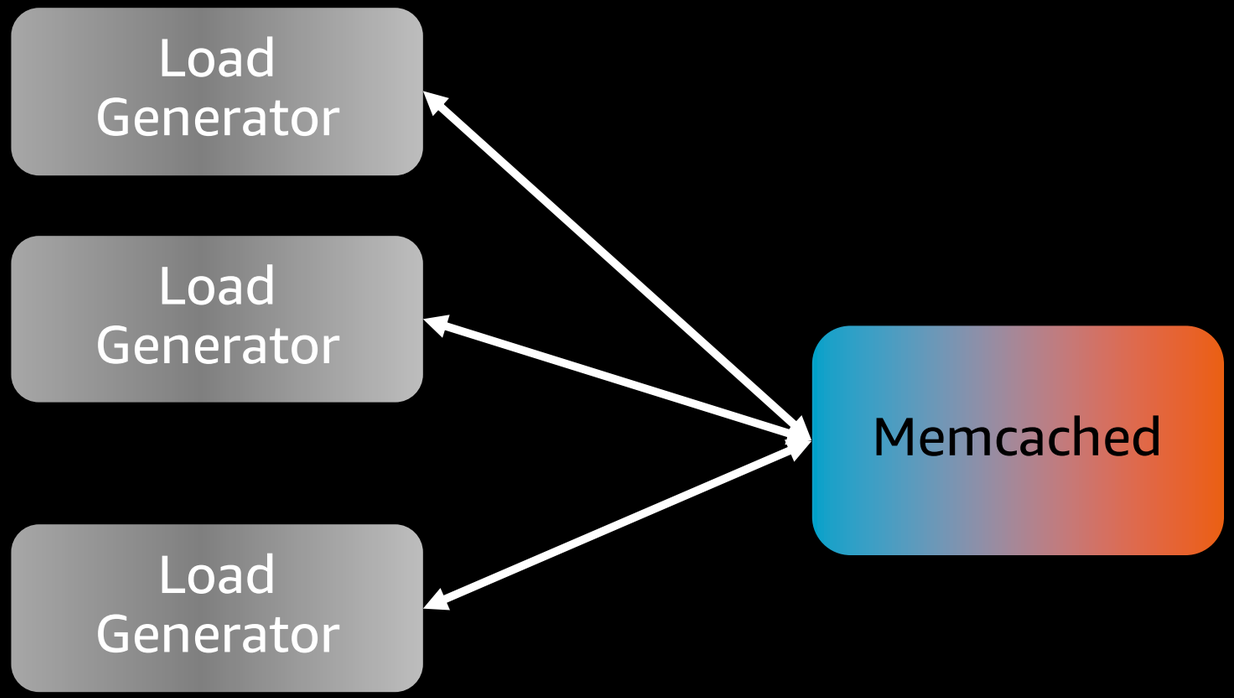

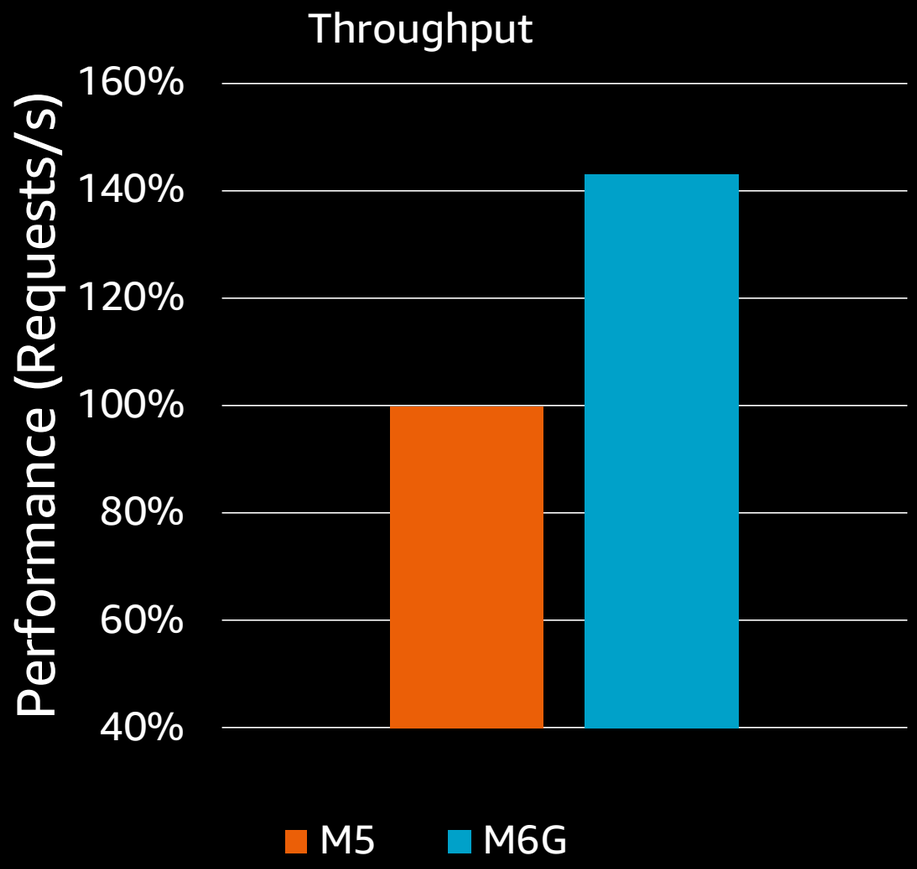

Memcached v1.5.16, 16B keys, 128B values, 7.8M KV-pairs, 576 connections for load generation from 2x c5.9xlarge instances, 16 additional connections used to measure latency from 1 additional c5.9xlarge,;each connection maintains 4096 outstanding requests; All servers in a cluster placement group

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

Memcached

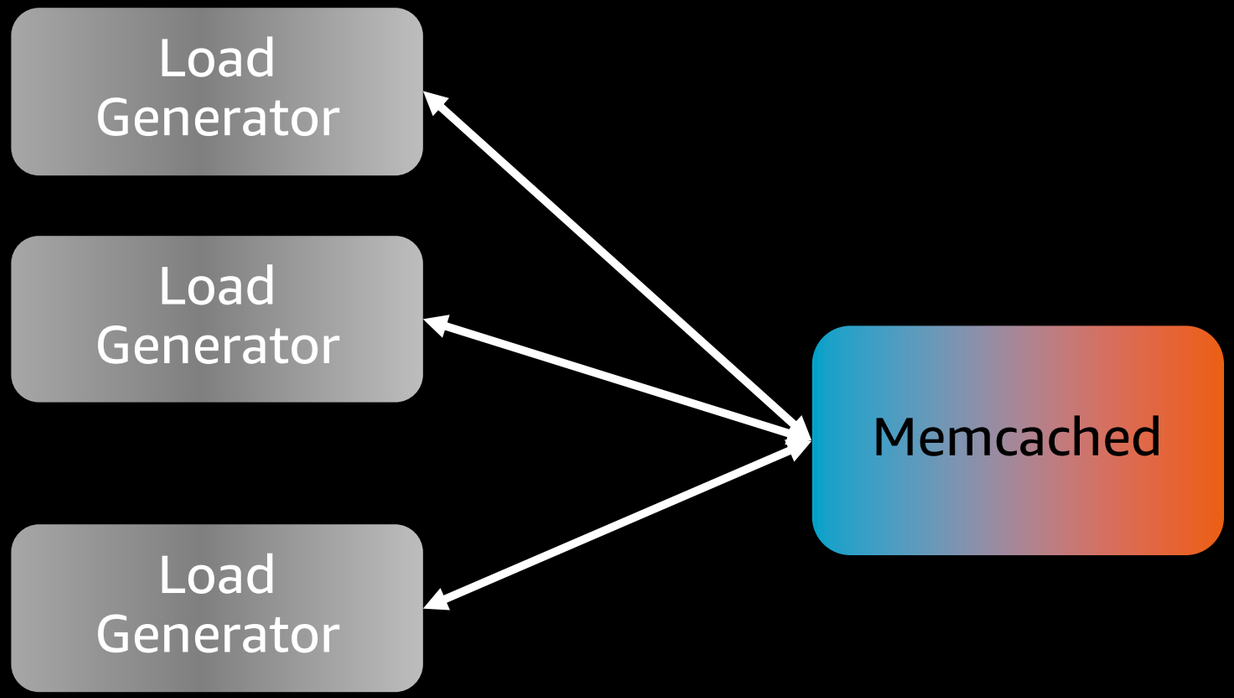

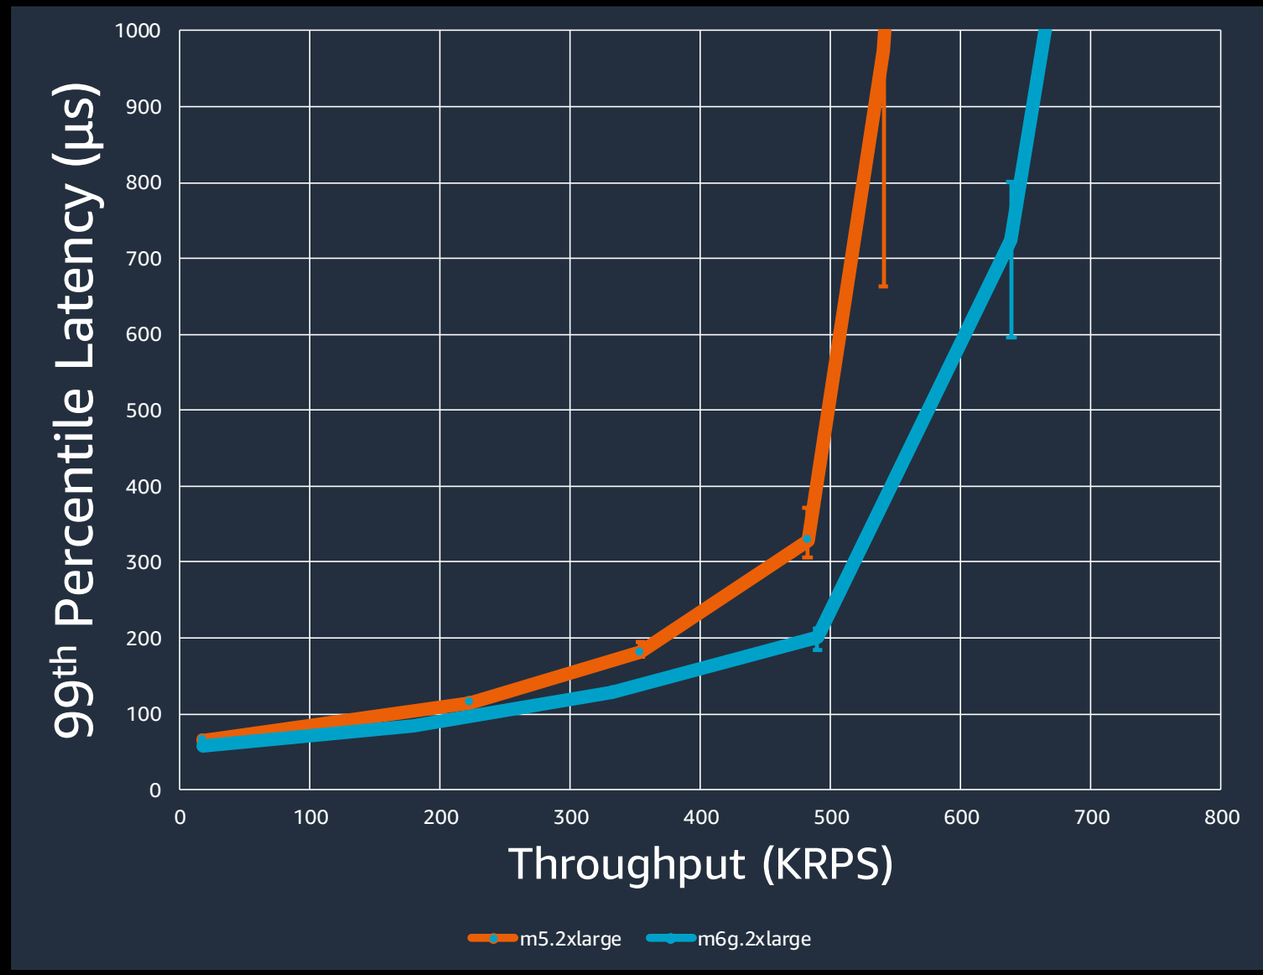

Memcached v1.5.16, 16B keys, 128B values, 7.8M KV-pairs, 576 connections for load generation from 2x c5.9xlarge instances, 16 additional connections used to measure latency from 1 additional c5.9xlarge, each connection maintains 4 outstanding requests; all servers in a cluster placement group

Media Encoding with x264

- Huge amount of video created daily
- Encoding it reduces bandwidth to deliver and storage of that video
- Using libx264 encoder encoded uncompressed 1080p to h264

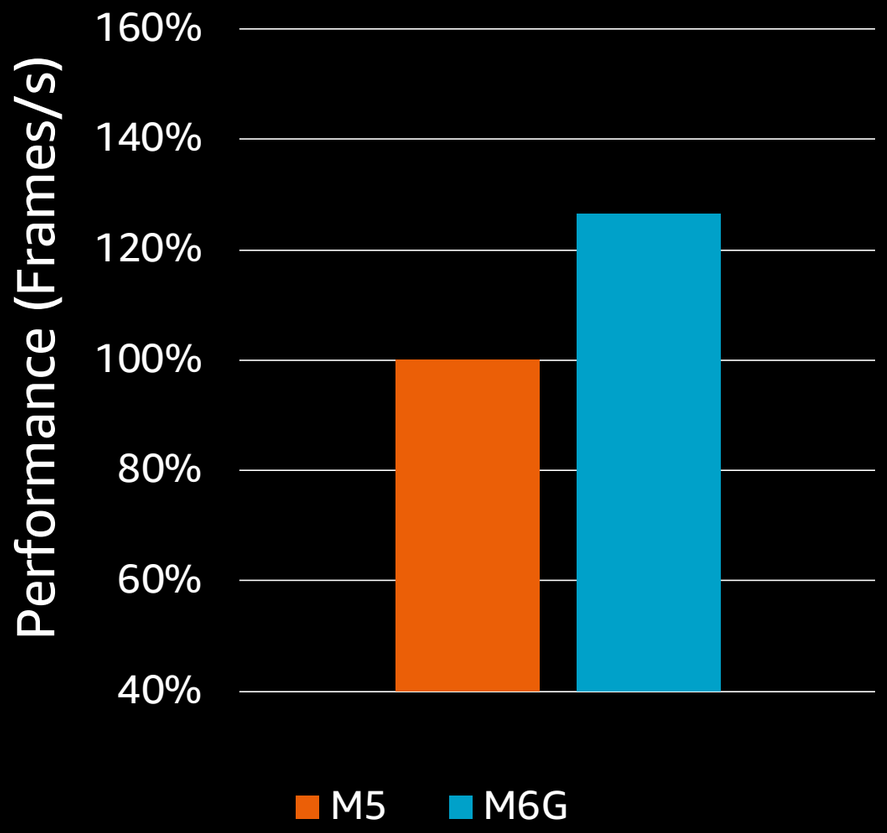

X264 (videolan.org version 3759fcb7), 4xl instance size, medium preset, input uncompressed 1080p50, output encoded h264 1080p50

Machine Learning

- BERT: The state-of-the-art text encoder
  - Feature representation (encoding) is crucial for ML
  - Deep neural network for text feature representation
- Graviton2 has fp16 and int8 support to accelerate Machine Learning workloads
- M6g can outperform M5
  - M5 with AVX-512 is limited to FP32
  - With FP16 support M6g performs better for CPU based inference

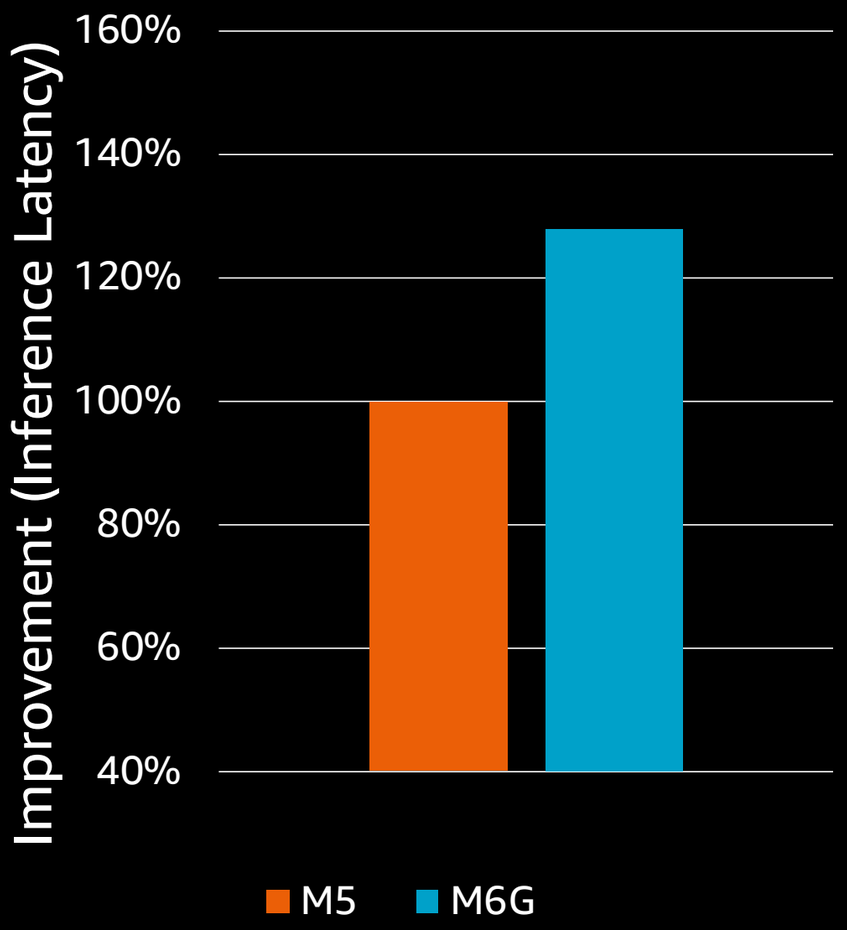

BERT classification using TVM and 64 length sequence on CPUs
Batch size of one; dedicated instances on xlarge size

EDA Performance – Arm and Cadence Xcellium

- Chip development is expensive
  - Design needs to be right the first time
- Arm uses millions of hours of CPU time per month to simulate their processor designs
  - Demand is highly variable depending on phase of the project
  - Perfect use-case for Amazon EC2
- Simulated Arm Cortex-A53 using Cadence Xcellium
  - 570 validation simulations on the the DPU RTL of the processor

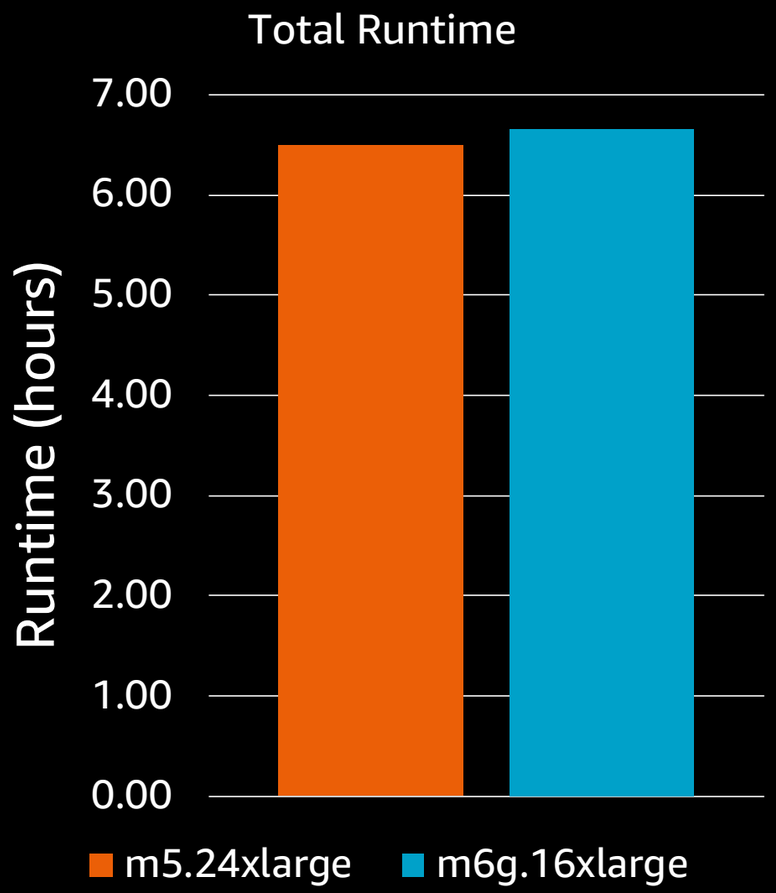

EDA Performance – Arm and Cadence Xcellium

- Chip development is expensive
  - Design needs to be right the first time
- Arm uses millions of hours of CPU time per month to simulate their processor designs
  - Demand is highly variable depending on phase of the project
  - Perfect use-case for Amazon EC2
- Simulated Arm Cortex-A53 using Cadence Xcellium
  - 570 validation simulations on the the DPU RTL of the processor

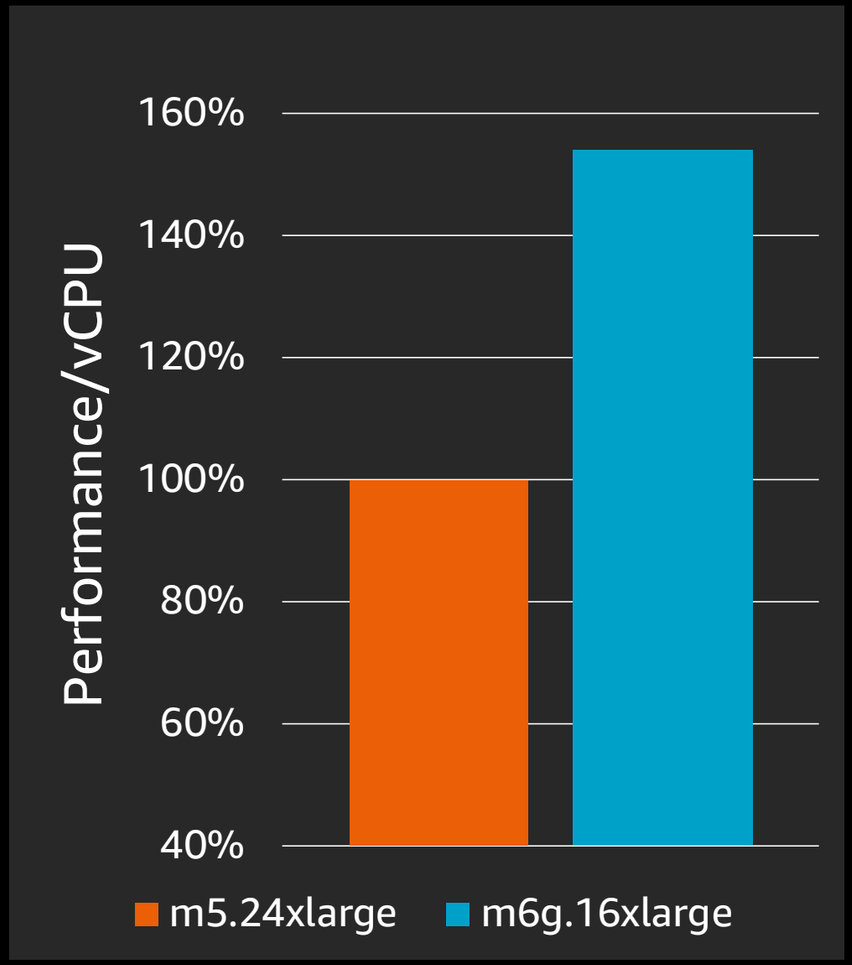

AWS Graviton delivers significantly better performance per vCPU for EDA

How about a database workload?

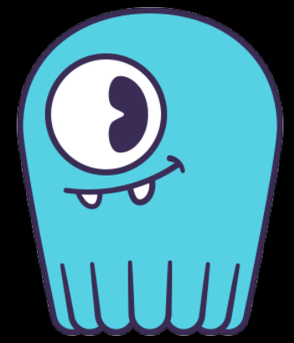

SCYLLA.

- Scylla is a high-throughput, low latency Big Data database
  - Thread-per-core architecture guarantees full CPU utilization
  - I/O Scheduler guarantees peak I/O throughput
  - Can easily reach AWS I3's 15GB/s bandwidth limit
- Can you use Scylla on A1 instance?
  - Yes. It works, and it is supported
  - But doesn't reach peak performance
- M6g instances change the game
  - All CPU and memory-bound workloads are supported
  - 4xlarge → 37.5k reads/s/core; 5x improvement over A1!
  - 64 vCPU theoretical limit around 2.4M reads/s
  - Amazon EBS is still instance storage

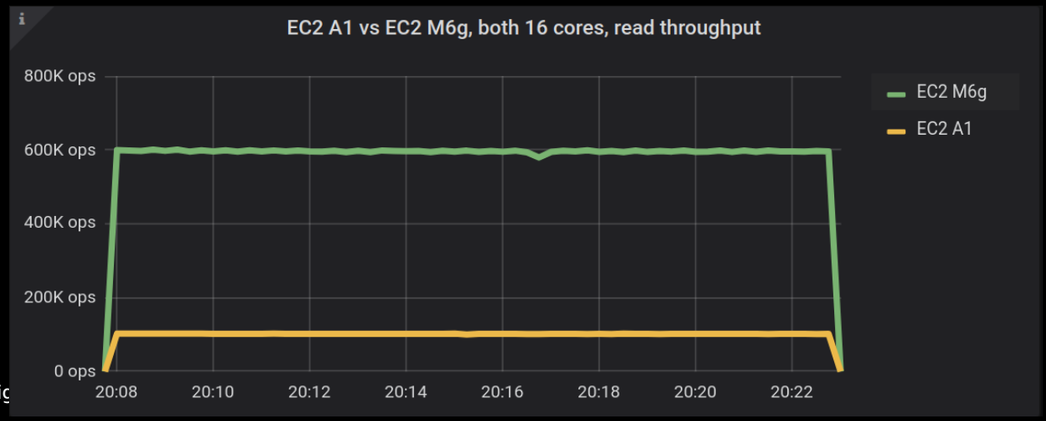

M6gd completes the story

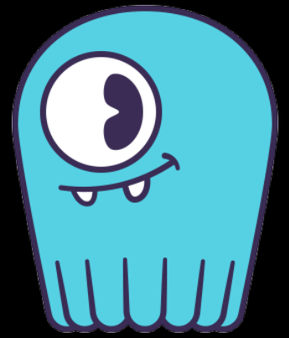

SCYLLA.

With a preview version of M6gd

- Over 5GB/s of disk bandwidth
- Close to 1M IOPS
- Enough for demanding storage workloads

Scylla welcomes the M6g series: ready for NoSQL

Graviton 2 实例拥有更低的总拥有成本

ARM 家族卓越性能

低于 M5系列 20% 成本

提供 40% 更高性价比

ARM 家族卓越体现

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营

感谢参加 AWS INNOVATE 2020 在线技术大会

我们希望您在这里找到感兴趣的内容！

也请帮助我们完成**投票打分**和**反馈问卷**。

欲获取关于 AWS 的更多信息和技术内容，可以通过以下方式找到我们：

微信订阅号：AWS 云计算（awschina）

微信服务号：AWS Builder 俱乐部（amazonaws）

新浪微博：https://www.weibo.com/amazonaws/

抖音：亚马逊云计算（抖音号：266052872）

视频中心：http://aws.amazon.bokecc.com/

博客：https://aws.amazon.com/cn/blogs/china/

更多线上活动 ： https://aws.amazon.com/cn/about-aws/events/webinar/

AWS 中国账户注册

AWS 全球账户注册

AWS INNOVATE ONLINE CONFERENCE

AWS 中国（宁夏）区域由西云数据运营
AWS 中国（北京）区域由光环新网运营